

HYPERSPECTRAL FEATURE DETECTION ONBOARD THE EARTH OBSERVING ONE SPACECRAFT USING SUPERPIXEL SEGMENTATION AND ENDMEMBER EXTRACTION

David R. Thompson⁽¹⁾, Benjamin Bornstein⁽¹⁾, Brian D. Bue⁽²⁾, Daniel Q. Tran⁽¹⁾, Steve A. Chien⁽¹⁾, Rebecca Castaño⁽¹⁾

⁽¹⁾ Jet Propulsion Laboratory, California Institute of Technology, Pasadena CA 91109 USA,
Email: firstname.lastname@jpl.nasa.gov

⁽²⁾ Rice University, 6100 Main St., Houston TX 77005 USA, Email: Brian.D.Bue@rice.edu

ABSTRACT

We present a demonstration of onboard hyperspectral image processing with the potential to reduce mission downlink requirements. The system detects spectral endmembers and then uses them to map units of surface material. This summarizes the content of the scene, reveals spectral anomalies warranting fast response, and reduces data volume by two orders of magnitude. We have integrated this system into the Autonomous Sciencecraft Experiment [1] for operational use onboard the Earth Observing One (EO-1) spacecraft. The system does not require prior knowledge about spectra of interest. We report on a series of trial overflights in which identical spacecraft commands are effective for autonomous spectral discovery and mapping for varied target features, scenes and imaging conditions.

1. INTRODUCTION

To date, the predominant approach for robotic space exploration has relied on instrument platforms executing scripted command sequences with data acquisition and downlink plans arranged long in advance. This is appropriate for static lander missions, or for mapping missions whose goal is a comprehensive survey. However, future generations of exploration spacecraft may face increasingly dynamic and complex environments that force new innovations in operations procedure [2]. Examples include: rover platforms travelling long distances [3], such as traverses already performed by the Mars Exploration Rovers [4]; and missions to primitive bodies, such as comet rendezvous or potential flyby tours of the Trojan objects. Only a small fraction of the potential targets in these environments would ever be measured, and plans could be revised during the mission. Here science data could play a *tactical* role to improve decisions about activities to perform in the next command cycle.

These missions could benefit from onboard data analysis in several ways [5]. First, onboard analysis could summarize the most bandwidth-expensive data products and increase the effective information content of each communication cycle [5]. This could tell operators which data products are worth downlinking

from the spacecraft, and ensure that key targets were not missed. Second, in the most dynamic environments, operators might gain additional advantage by shifting time-critical decisions across the light time delay so that the spacecraft itself recognizes transient targets of opportunity and triggers needed response actions such as followup measurements [1]. Such needs are greatest in situations when round-trip communication is impossible, such as flybys or short-lived landers.

Onboard analysis holds special benefit for imaging spectrometers (alternatively, hyperspectral imagers). Imaging spectrometers typically provide measurements in Visible Near Infrared (VNIR) or Infrared (IR) ranges, with up to hundreds of bands at spectral resolutions down to 10-40nm or better [6,7]. These data combine wide area coverage with compositional insight, giving invaluable information for physical modelling and interpretation. Not surprisingly, a significant fraction of deep space exploration missions employ them. However, due to the large number of bands since a single scene may demand several hundred megabytes of downlink bandwidth. This size is a serious disadvantage (for example, outweighing a typical daily bandwidth allowance for the Mars Exploration Rover extended mission), and as a result imaging spectrometers are generally limited to a very small fraction of their potential duty cycles. Onboard data analysis may help ameliorate this issue [8]. By autonomously identifying spectral anomalies and other targets of interest, the remote explorer can discover features that warrant followup measurements. It can draft maps of surface composition, enabling bandwidth-efficient summaries where full maps are infeasible.

Automating hyperspectral analysis for spacecraft is a difficult problem due to unknown spectral compositions of the target. Unknown spectral shapes make the targets difficult to specify in advance, so planetary spectra typically reveal their secrets only after rigorous post-analyses and calibration by trained analysts. Naturally, this is not possible for autonomous operations that must recognize unanticipated features in the scene and function correctly on the first try. Another challenge is computational efficiency, since computational resources

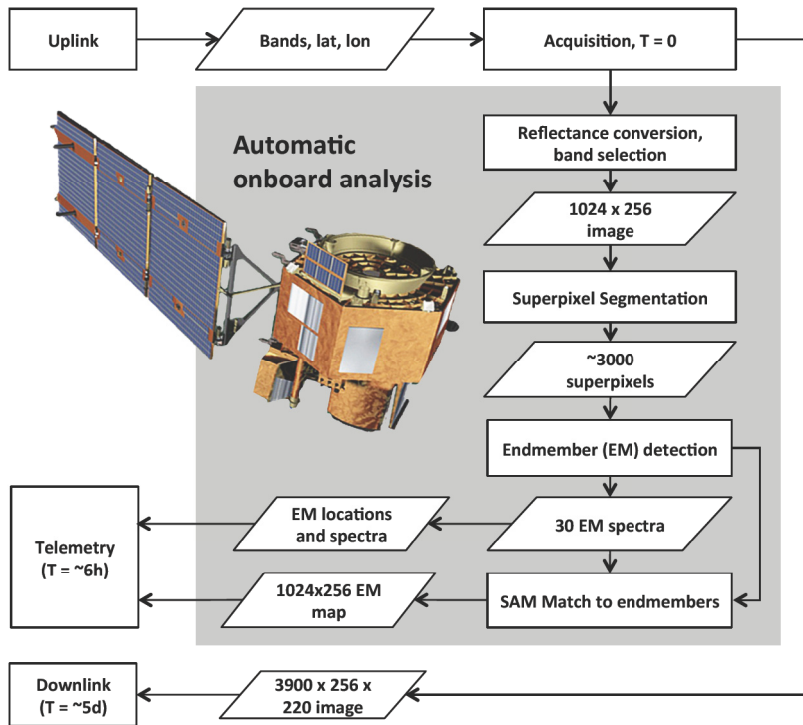


Figure 1: Procedure for onboard processing and summary downlink by EO-1. The procedure has three steps. In step (1) preprocessing, the image is reduced to a 11-band subset to satisfy EO-1 software constraints. Step (2), superpixel segmentation, transforms the image into a new representation based on spectrally-homogeneous contiguous areas. This improves noise while reducing the data set size by a factor of ~100. In step (3) endmember detection, numerical methods find numerically extreme superpixels corresponding to the spectral endmembers. The automatic feature detection acts as a “shortcut” to the standard downlink, permitting a summary product of 50KB to be delivered within 6 hours of acquisition. **EO-1 Image: NASA/GSFC**

aboard spacecraft are highly limited. Systems using Field Programmable Gate Arrays (FPGAs) have been developed for spaceflight applications [8], but none have yet been deployed to orbit. A final complication is a noise level that is generally much worse than airborne applications due to thermal variation and unconstrained imaging conditions. For these reasons, proofs-of-concept are needed to demonstrate that meaningful hyperspectral analysis can indeed be performed in a flight environment.

This work reports initial results from a prototype system for autonomous hyperspectral summary and mapping, which was deployed to the Earth Observing One (EO-1) spacecraft [9] in the fall of 2011. It has been in operational use since this time, performing multiple overflight tests over a variety of scenes with mineralogical and astrobiology interest. Our approach uses a superpixel endmember detection method [10] to identify the most salient spectra in each scene. An initial segmentation provides resilience to noise while reducing computation requirements of subsequent processing. Then, an endmember detection procedure finds the spectra representing the strongest distinctive spectral signatures. This reliably retrieves both representative and anomalous features. The spacecraft then computes the location and extent of these spectral signatures, forming a simple but efficient summary map. The following section describes the basic analysis procedure, detailed previously in Bornstein et al [11]. We then describe initial demonstration results, and close with a discussion of performance and the implications of these trials for future technology development and

mission use.

2. APPROACH

The Hyperion instrument aboard EO-1 is a hyperspectral imager with 220 bands ranging from Visible to Short Wave Infrared (SWIR) frequencies. During its extended mission it has hosted a range of autonomous data analysis software for detecting features such as floods and volcanism through the Autonomous Sciencecraft Experiment (ASE) [12-15]. We have integrated spectral endmember detection and mapping into this system. The basic procedure appears in Figure 1. The normal EO-1 image downlink protocol can transmit a full image within a few days. In contrast, the onboard system creates a summary product less than 50KB in size, which fits into a direct downlink within 6 hours of acquisition.

2.1. Preprocessing

The first step of the onboard processing transforms the raw radiance data into a *pseudo-reflectance* product. This computation is based on distance to the sun, solar zenith, and the solar input in each spectral channel. The method is identical to onboard reflectance computation used in other ASE experiments and described at greater length in [16]. During this conversion the preprocessing selects 12 channels from the 220 Hyperion instrument spectral bands. This satisfies constraints on working memory based on the flight avionics system and the intrinsic hardware capabilities of the EO-1 processor. An additional constraint is that one band from each of two detectors must be loaded. The Hyperion instrument

has two different detector elements covering the visible and infrared regions. In practice, the reflectance values the detectors can differ significantly. This could cause scaling problems for purely data-driven analyses that used channels from both detectors simultaneously. As compensation, we limit our analysis to one of the two detectors. After satisfying the constraint that both detectors be queried this leaves a free allowance of 11 channels in either spectral region that are available for use onboard. The band selection can change with each acquisition if required. All of the experiments in this work used one of two band set options that appear in Table 1. The subsets are a compromise that provides even spectral coverage over organic and mineral absorption features, while avoiding specific channels that appeared particularly noisy in downlinked images.

SWIR Band	nm	VNIR Band	nm
196	2113.0	16	508.2
198	2133.2	20	548.9
202	2173.5	25	599.8
204	2193.7	30	650.7
205	2203.8	33	681.2
207	2224.0	36	711.7
208	2234.1	39	742.3
211	2264.3	42	772.8
214	2294.6	45	803.3
219	2345.1	48	833.8
221	2365.2	52	874.5

Table 1: Channels used for the Short Wave Infrared (SWIR) and Visible Near Infrared (VNIR) band sets.

The preprocessing operates on a subframe of the typical Hyperion swath that is 1024x256 pixels in size. It is centered to within spacecraft pointing accuracy on the overflight latitude/longitude location.

2.2. Superpixel Segmentation

The preprocessed data product is not yet suited for unsupervised endmember detection. It exhibits artifacts such as column noise and edge effects that are typical of orbital hyperspectral remote sensing. Individual pixels may also have high levels of shot noise and intrinsic white measurement noise. Moreover, the image subframe containing over 256K pixels is still too large for an efficient endmember analysis: such algorithms typically scale polynomially, and are too intensive for the 20MHz EO-1 flight computer. Even simpler operations such as image compression typically take up to an hour for a standard Hyperion image.

These considerations motivate a second intermediate step in which we transform the pixelwise representation to a set of spatially contiguous segments. Superpixel representations [10] oversegment the image, shattering it into spectrally homogeneous pixel clusters. It is natural to interpret these as physical features (or parts of

physical features) in the scene, suggesting the additional constraint that segments be spatially contiguous. This makes computing them a fast operation. Here we use an agglomerative segmentation algorithm [17] based on disjoint set merging, which scales favorably by $n \log n$ for n pixels. It can accommodate any distance measure; here we use Euclidean distance. Recent studies suggest that learned metrics can achieve even better results by learning task-specific notions of spectral similarity [18]. After identifying superpixel areas, the onboard system represents each segment using its average spectrum. Assuming that each the superpixel lies entirely on a physically homogeneous feature, the underlying reflectance is the same in all pixels and averaging achieves a noise reduction according to the square root of the superpixel size [10]. This new representation also reduces the dataset size by multiple orders of magnitude, uniquely enabling the endmember detection procedure that follows.

2.3. Endmember Detection

Spectra observed in remote sensing data are often assumed to obey an *aerial mixing* assumption. Here the scene is comprised of a small handful of *endmember* materials [19]. The endmember reflectances combine linearly in proportion to their area of extent within the ground footprint to generate the measured reflectance of the image pixels. We make the second simplifying assumption that the “purest” examples of each endmember appear somewhere in the scene. In other words, we assume that we could reconstruct all the spectra in the data set by linear combination of these pure spectra if appropriate mixing ratios were known.

The third stage of EO-1 processing uses a Sequential Maximum Angle Convex Cone (SMACC) strategy to find the endmembers [20]. This method proceeds by expanding a basis set of endmembers, iteratively maximizing the volume of their associated convex cone. Intuitively, the extreme points in this convex region are linear endmembers, and the larger their volume the more distinct spectra their linear combinations can represent. Many other endmember detection strategies are possible, and we refer the reader to the taxonomy of [21]. We find SMACC effective for this system because it is efficient, deterministic and reproducible. Moreover, it is a sequential approach that can return any number of endmembers in ranked order. This contrasts with “simultaneous” algorithms that fit all endmembers at once and require that the user know the number of endmembers in advance.

In this work, we typically compute the top 30 endmembers. This is far larger than the intrinsic dimensionality of the Hyperion data, which is typically less than 10. Extra endmembers provide a margin of error and allow for noise and outlier features that are not

compositionally meaningful. Some noise is unavoidable due to clouds, shadows, and other high-variance areas. Despite the superpixel representation's noise reduction, some of these features inevitably appear as numerically-extreme data points that make their way into the endmember list. Moreover, the spectra of a scene do not always obey a pure-pixel linear mixing assumption, and compositionally redundant endmembers may also be selected. Computing an overcomplete set of 30 endmembers provides a margin of error so that unique physically meaningful features are not crowded out by these other artifacts. It is typically straightforward to identify the important and redundant spectra on the ground, since the onboard system has already performed the essential work of finding the basis spectra and their locations in the scene. The extra endmembers impose little bandwidth or computation cost, and the bandwidth needed by their spectra is negligible relative to the other summary products.

Figure 2 shows an example segmentation based on an actual overflight of Blood Falls, Antarctica. This scene contains a hidden feature of interest: bright red iron oxidizes visible in a saline outflow from a subglacial system near Taylor Glacier [22]. Panel A shows the context image, while panel B shows the superpixel representation. Here homogenous regions of land and ice have been grouped together into clusters. This facilitates the efficient detection of endmember spectra. The region of the outflow appears as a distinctive spectral endmember, identified by the superpixel indicated in panel C. This was produced based on the map that the system generated onboard and downlinked within hours. Panel D shows an image from the surface.

2.4. Spectral Angle Mapping

After endmember detection, the system computes a spectral angle match between each superpixel and the entire set of endmembers. The spectral angle is simply a normalized dot product measuring the similarity between each spectrum and each endmember. It produces a classification map relating each pixel of the original image to the best-matching endmember spectrum. The final downlink includes multiple summary products:

1. The complete list of endmember spectra, with their pseudo-reflectance values in each of the 11 selected bands,
2. The pixel locations of these endmember superpixels in the scene, identifying the locations of the salient features,
3. An integer map relating every pixel in the image to the endmember spectrum that matches it best. It serves as a simple "draft" compositional map.

These products have value as an efficient summary, but could also enable onboard decision-making to respond to unexpected spectra. For example, the system could be




Figure 2: Blood falls overflight illustrating the detection process. (A) Original 1024x256 pixel pseudoreflectance image (here, three visible RGB channels are shown). (B) Superpixel segmentation identifies contiguous spectrally-homogeneous areas. (C) The superpixel corresponding to a salient endmember spectrum at the outflow location. (D) The blood falls feature as viewed from the ground Image: United States Antarctic Program.

Scene	Bands	Detections
1. Cuprite	SWIR	Alunite, Muscovite, Calcite
2. Cuprite	SWIR	Alunite, Muscovite, Calcite, Kaolinite
3. Steamboat	SWIR	Silica, Alunite/Kaolinite
4. Mammoth	VNIR	Thermal springs []
5. Mammoth	VNIR	None (70% cloud cover) []
6. Blood Falls	VNIR	Glacial outflow

Table 2: Detection results from EO-1 test overflights.

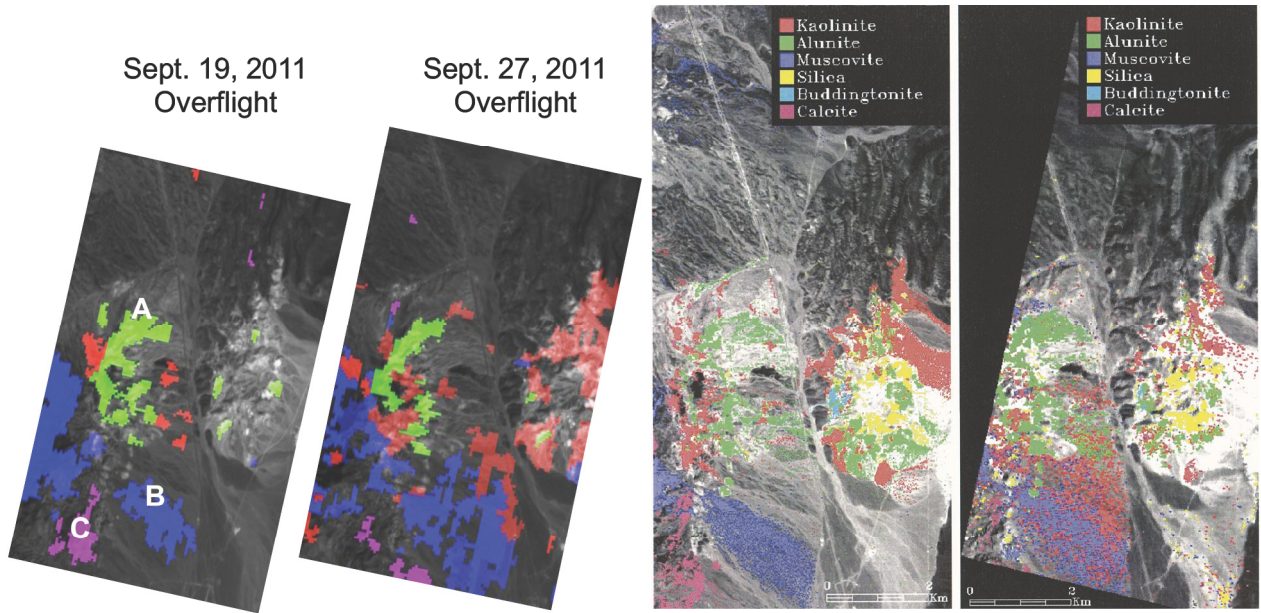


Figure 5: Two overflights of Cuprite, Nevada. We show products based on the fully-autonomous onboard detection and mapping (Extreme and Center Left). Detections include Alunite (A), Muscovite (B), and Calcite (C). The red area contains a mixture with features of Kaolinite and Alunite spectra. We color the image to show the correspondence with manually-produced maps of the same region based on AVIRIS (Center Right) and Hyperion (Extreme Right). Map image courtesy Kruse et al. [24].

instructed to return a full image on detection of a specific target endmember, or upon the appearance of a wholly novel spectral signature.

3. EXPERIMENTAL RESULTS

We commanded test overflights during the Fall of 2011 over several sites of mineralogical interest, including well-studied locales where extensive ground truth data was available. These include: Cuprite, a mining district in Nevada, which contains a range of minerals with SWIR absorption features [23,24]; Steamboat Springs, an active thermal system with a terrace and exposed SWIR mineral signatures [25]; Mammoth Hot Springs, an active thermal spring system in Yellowstone National Park, USA [26]; and Blood Falls, the saline subglacial outflow appearing in Figure 2 [22]. *Apart from the overflight target latitude and longitude, and the choice of either the VNIR or SWIR band set, we held all parameters of the detection system constant across all trials.* Table 2 shows the complete list of runs. All resulted in successful detections, except for the second overflight of Mammoth Hot Springs. In that case there was no endmember corresponding to the hot springs feature; the system was confused by overwhelming cloud cover (over 70% of the scene, according to the standard Hyperion product estimate) so these features dominated the endmember list.

Figure 5 shows a comparison of the two Cuprite overflights with canonical maps produced by Kruse et al. of the same scene using AVIRIS and Hyperion data




Figure 4: Endmember spectral features appearing in the Cuprite Scene. Letters correspond to the map regions above. All endmember spectra are interpretable as having diagnostic spectral shapes corresponding to the target mineral classes. Alunite contains an additional absorption feature near 2320nm which is not evident in this spectrum do to sparse spectral sampling. However, the feature is visible in the full-spectrum downlink and confirms the interpretation.

[24]. We select colors to highlight the correspondences between the different mapping runs. We note that there is some intrinsic variability in the detection result based on factors such as Signal to Noise, solar zenith, and atmospheric conditions. Many of the same features are discovered by the onboard system, including Alunite, Muscovite, and Calcite spectra. The downlinked spectra associated with these regions appear in Figure 4.

Figure 5 shows the run over Steamboat Springs, Nevada, which also evidences good agreement with formal mineralogical maps of Kruse et al. [24]. The two regions colored regions show the extent of Silica and Alunite/Kaolinite features that were automatically detected onboard. As in the Cuprite case, the manual maps were produced by extensive ground-truth analysis and an expert understanding of the site.

4. DISCUSSION AND CONCLUSIONS

Initial results of the EO-1 system are encouraging; to date, the system has proven effective for both target detection and scene summary, for both dynamic and mineralogical features of interest. To our knowledge, the work described here is the first instance of onboard spectral discovery in which the spacecraft detects outliers that were not anticipated in advance.

While the system is limited by intrinsic hardware and software constraints, it offers an initial proof-of-concept for the techniques and a starting point for future progress. It is likely that future systems will use some combination of lossy compression, endmember detection, and matched filtering [27]. Such systems may be implemented in more specialized hardware such as FPGAs or multi-core architectures. Advances in data throughput will permit analysis of more spectral bands, improving total SNR and mapping performance. The methods we describe here may also apply to multiband framing cameras such as the MER Pancam or the OSIRIS instrument aboard the Rosetta Spacecraft. The dual ability to generate summary maps and to identify targets of opportunity for followup measurements can be a powerful tool to surmount communications limits and facilitate agile remote exploration.

5. ACKNOWLEDGEMENTS

We acknowledge the assistance of the Autonomous Sciencecraft Experiment, the Earth Observing One mission the GSFC Science team including Elizabeth Middleton, Stephen Ungar, Petya Campbell, and Lawrence Ong. We also acknowledge the operations support from Dan Mandl and Stuart Frye. We acknowledge support from a Technology Development Grant under the Advanced Multimission Operating System (AMMOS) and the Multimission Ground Support Services (MGSS) office. We thank Lukas Mandrake and Martha Gilmore for their help in




Figure 5: Steamboat Springs, Nevada. The thermal springs evidences Silica (A) and Alunite/Kaolinite (B), each of which appear as distinct endmembers with associated regions shown at left. The rightmost panel shows a result by Kruse et al. based on their manual analysis of a different overflight using Mixture Tuned Matched Filter (MTMF) mapping [25].

developing the superpixel segmentation approach. We also thank Fred Kruse for his assistance and the use of previous work. Finally, we thank Aaron Noell and Damhnait Gleeson for their support, counsel, and encouragement. The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology. Copyright 2012 California Institute of Technology. All Rights Reserved. U.S. government support acknowledged.

6. REFERENCES

1. Sherwood, D. Tran, B. Cichy, G. Rabideau, R. Castano, A.G. Davies, D. Mandl, S. Frye, B. Trout, S. Shulman, and D. Boyer. (2005). Using autonomy flight software to improve science return on Earth Observing One. *Journal of Aerospace Computing, Information, and Communication*, April 2005.
2. Thompson, D. R., J. Castillo-Rogez, S. A. Chien, R. Doyle, T. Estlin, and D. McLaren (2012). Agile science operations: A new approach for primitive bodies exploration. *Proceedings of the SpaceOps Conference, Stockholm, Sweden*.
3. Thompson, D. R., D. Wettergreen, and F. C. Peralta, (2011). Autonomous science during large-scale robotic survey. *Journal of Field Robotics*, 28 (4), p. 542–564.
4. Estlin, T. B. Bornstein, D. Gaines, R. C. Anderson, D. R. Thompson, M. Burl, R. Castano, and M. Judd (2012). AEGIS automated targeting for the MER opportunity rover. *ACM Trans. Intelligent Systems Technology* 3(3).
5. R.Castano, T.Estlin, R.Anderson, D.Gaines, A.Castano, B.Bornstein, C. Chouinard, and M.Judd (2007) OASIS: Onboard autonomous science nvestigation system for opportunistic rover science. *Journal of Field Robotics*, 24 (5), p. 379–

6. S. Murchie et al., (2007) Compact Reconnaissance Imaging Spectrometer for Mars (CRISM) on Mars Reconnaissance Orbiter. *J. Geophys. Res.*, 112 (5).
7. R. O. Green, et al. (2011). The Moon Mineralogy Mapper (M3) imaging spectrometer for lunar science: Instrument description, calibration, on-orbit measurements, science data calibration and on-orbit validation. *J. Geophys. Res.*, doi: 10.1029/2011JE003797.
8. S. Qian, M. Bergcron, I. Cunningham, L. Gagnon, and A. Hollinger (2008), Near lossless data compression onboard a hyperspectral satellite, *IEEE Trans. Aerospace and Electronic Systems*, 42 (3), pp. 851–866.
9. S. Ungar, J. Pearlman, J. Mendenhall, and D. Reuter, “Overview of the Earth Observing One (EO-1) mission,” *IEEE Trans. Geosci. Remote Sens.* 41 (6), p. 11491159, 2003.
10. Thompson, D. R., L. Mandrake, M.S. Gilmore, and R. Castaño (2010). Superpixel endmember detection. *IEEE Transactions on Geoscience and Remote Sensing*. 48 (11), pp. 4023–4033.
11. Bornstein, B., D. R. Thompson, D. Tran, B. Bue, S. Chien, R. Castano (2011). Efficient spectral endmember detection onboard the EO-1 spacecraft. *IEEE WHISPERS*.
12. S. Chien, R. Sherwood, D. Tran, B. Cichy, G. Rabideau, R. Castano, A. Davies, D. Mandl, S. Frye, B. Trout, S. Shulman, and D. Boyer (2005). Using autonomy flight software to improve science return on Earth Observing One, *Journal of Aerospace Computing, Information, and Communication*.
13. F. Ip, J. Dohm, V. Baker, T. Doggett, A. G. Davies, R. Castano, S. Chien, B. Cichy, R. Greeley, and R. Sherwood (2006) Development and testing of the Autonomous Spacecraft Experiment (ASE) floodwater classifiers: real-time smart reconnaissance of transient flooding, *Remote Sensing of Environment*, 101 (4), pp. 463–481.
12. A. G. Davies and S. Chien and V. Baker and T. Doggett and J. Dohm and R. Greeley and F. Ip and R. Castano and B. Cichy and R. Lee and G. Rabideau and D. Tran and R. Sherwood, “Monitoring active volcanism with the Autonomous Spacecraft Experiment (ASE),” *Remote Sensing of Environment*, vol. 101, no. 4, pp. 427–446, 2006.
15. Doggett, T., R. Greeley, A. G. Davies, S. Chien, B. Cichy, R. Castano, K. Williams, V. Baker, J. Dohm, and F. Ip, Autonomous on-board detection of cryospheric change, *Remote Sensing of Environment*, vol. 101, no. 4, pp. 447–462, 2006.
16. K. R. Knapp (1996), Radiance to reflectance for GOES-8 channel 1, *CO St. University Tech Rept.*
17. Felzenszwalb, P.F., D.P. Huttenlocher (2004), Efficient graph-based image segmentation. *International Journal of Computer Vision*, 49 (2).
18. Bue, B., Thompson, D. R., Gilmore, M. S., Castano, R. (2011). Metric learning for hyperspectral image segmentation., *IEEE WHISPERS*.
19. Keshava, N. and J. Mustard, (2002) Spectral unmixing. *IEEE Signal Processing Magazine*, 19 (1), pp. 44–57.
20. Gruninger, J., A. Ratkowski, and M.L. Hoke (2004). The sequential maximum angle convex cone (SMACC) endmember model. *Proceedings SPIE, Algorithms for multispectral, hyperspectral and ultraspectral Imagery*. 5429. p.1-14.
21. A. Plaza, P. Martinez, R. Pérez, and J. Plaza, (2006). A quantitative and comparative analysis of endmember extraction algorithms from hyperspectral data. *IEEE Trans. Geoscience and Remote Sensing*, 42 (3), pp. 650–663.
22. J. Mikucki, A. Pearson, D. Johnston, A. Turchyn, J. Farquhar, D. Schrag, A. Anbar, J. Priscu, and P. Lee (2009) A contemporary microbially maintained subglacial ferrous “ocean.” *Science*, 324 (5925), p. 397–400.
23. G. A. Swayze (1997) The hydrothermal and structural history of the Cuprite Mining District, Southwestern Nevada: An integrated geological and geophysical approach. *Ph.D. thesis, Univ. Colorado, Boulder, CO*.
24. F. A. Kruse, J. W. Boardman, and J. F. Huntington (2003). Comparison of airborne hyperspectral data and EO-1 Hyperion for mineral mapping. *IEEE Trans. Geosci. Remote Sens.*, 41 pp. 1388–1400.
25. F. A. Kruse (1999). Mapping Hot spring deposits with AVIRIS at Steamboat Springs, Nevada. *Proceedings of the 8th JPL Airborne Earth Science Workshop, Jet Propulsion Laboratory Publication*, 99 (17), pp. 239 – 246.
26. K. Livo, F. A. Kruse, R. Clark, R. Kokaly, and W. I. Shanks (2007). Hydrothermally altered rock and hot-spring deposits at Yellowstone National Park characterized using airborne visible- and infrared-spectroscopy Data. *U.S. Geological Survey Professional Paper 1717*. Ch. 0.
27. A. Plaza, Q. Du, Y. Chang, and R. King (2011). High performance computing for hyperspectral remote sensing. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 4 (3), pp. 528–544.